


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: ☒ The ACM Digital Library ☐ The Guide

Searching within **The ACM Digital Library** with **Advanced Search**: ("b-number" and operator and prefix) (s new search)

Found 16 of 248,317

REFINE YOUR SEARCH

[Search Results](#)
[Related Journals](#)
[Related SIGs](#)
[Related Conferences](#)

▼ Refine by Keywords

[Discovered Terms](#)

▼ Refine by People

[Names](#)
[Institutions](#)
[Authors](#)
[Reviewers](#)

▼ Refine by Publications

[Publication Year](#)
[Publication Names](#)
[ACM Publications](#)
[All Publications](#)
[Publishers](#)

▼ Refine by Conferences

[Sponsors](#)
[Events](#)
[Proceeding Series](#)

Results 1 - 16 of 16

Sort by in

[Save results to a Binder](#)

1 [Rule-based multi-query optimization](#)

Mingsheng Hong, Mirek Riedewald, Christoph Koch, Johannes Gehrke, Alan Demers
 March 2009 **EDBT '09**: Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology

Publisher: ACM

Full text available: Pdf (731.50 KB) [Additional Information: full citation, abstract, references](#)

Bibliometrics: Downloads (6 Weeks): 26, Downloads (12 Months): 26, Citation Count: 0

Data stream management systems usually have to process many long-running queries that are active at the same time. Multiple queries can be evaluated more efficiently together than independently, because it is often possible to share state and computation.

2 [Tailor-made lock protocols and their DBMS integration](#)

Sebastian Bächle, Theo Härder

March 2008 **SETMDM '08**: Proceedings of the 2008 EDBT workshop on Software engineering for tailor-made data management

Publisher: ACM

Full text available: Pdf (451.98 KB) [Additional Information: full citation, abstract, references](#)

Bibliometrics: Downloads (6 Weeks): 4, Downloads (12 Months): 25, Citation Count: 0

We outline the use of fine-grained lock protocols as a concurrency control mechanism, the collaboration on XML documents and show that their tailor-made optimization towards the access model used (e.g., DOM operations) pays off. We discuss how he wired ...

3 [FMware: middleware for efficient filtering and matching of XML messages with data](#)

K. Selçuk Candan, Mehmet E. Dönderler, Yan Qi, Jaikannan Ramamoorthy, Jong W. Ki
 November 2006 **Middleware '06**: Proceedings of the ACM/IFIP/USENIX 2006 International Conference on Middleware

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (1.84 MB) [Additional Information: full citation, abstract, references](#)

Bibliometrics: Downloads (6 Weeks): 1, Downloads (12 Months): 1, Citation Count: 0

XML message filtering systems are used for sifting through real-time messages to support business data mining and reporting. An XML message filtering system needs (a) process registered filter predicates on multiple distributed real-time streams and

ADVANCED SEARCH

[Advanced Search](#)

FEEDBACK

[Please provide us with feedback](#)

Found 16 of 248,317

Keywords: XML messaging, cluster-domain hashing, matching with local data

4 [More dynamic imperative languages](#)



A. Savidis

December 2005 **SIPLAN Notices**, Volume 40 Issue 12

Publisher: ACM

Full text available: [Pdf](#) (346.24 KB) **Additional Information:** [full citation](#), [abstract](#), [references](#), [index ter](#)

Bibliometrics: Downloads (6 Weeks): 3, Downloads (12 Months): 31, Citation Count: 0

We pursue entirely dynamic constructs for imperative object-oriented programming languages by shifting responsibilities from compile-time handling to runtime management. This accounts to the full range of language elements, going beyond statically defined: ...

Keywords: dynamic languages, dynamically extensible semantics, polymorphic pa programming, runtime inheritance

5 [Towards robust semantic role labeling](#)

Sameer S. Pradhan, Wayne Ward, James H. Martin

June 2008 **Computational Linguistics**, Volume 34 Issue 2

Publisher: MIT Press

Full text available: [Pdf](#) (165.08 KB) **Additional Information:** [full citation](#), [abstract](#), [references](#), [index ter](#)

Bibliometrics: Downloads (6 Weeks): 12, Downloads (12 Months): 98, Citation Count: 0

Most semantic role labeling (SRL) research has been focused on training and evaluation on the same corpus. This strategy, although appropriate for initiating research, can lead to overtraining to the particular corpus. This article describes the operation ...

6 [Combining flexibility and scalability in a peer-to-peer publish/subscribe system](#)

Chi Zhang, Arvind Krishnamurthy, Randolph Y. Wang, Jaswinder Pal Singh

November 2005 **Middleware '05: Proceedings of the ACM/IFIP/USENIX 2005 International Conference on Middleware**

Publisher: Springer-Verlag New York, Inc.

Full text available: [Pdf](#) (448.11 KB) **Additional Information:** [full citation](#), [abstract](#), [references](#)

Bibliometrics: Downloads (6 Weeks): 3, Downloads (12 Months): 3, Citation Count: 0

The content-based publish/subscribe model has been adopted by many services to deliver data between distributed users based on application-specific semantics. Two issues in such systems, the semantic expressiveness of content matching and the scalability ...


Keywords: matching, peer-to-peer, publish/subscribe

7 [A 50-Gb/s IP router](#)

Craig Partridge, Philip E. Carvey, Ed Burgess, Isidro Castineyra, Tom Clarke, Lise Grah Michael Hathaway, Phil Herman, Allen King, Steve Kohalmi, Tracy Ma, John Mcallen, Ti Mendez, Walter C. Milliken, Ronald Pettyjohn, John Rokosz, Joshua Seeger, Michael Se Steve Storch, Benjamin Tober, Gregory D. Troxel

June 1998 **IEEE/ACM Transactions on Networking (TON)**, Volume 6 Issue 3

Publisher: IEEE Press

Full text available:  Pdf (133.28 KB) Additional Information: [full citation](#), [references](#), [cited by](#), [index terms](#), [review](#)

Bibliometrics: Downloads (6 Weeks): 7, Downloads (12 Months): 89, Citation Count: 39

Keywords: data communications, internetworking, packet switching, routing

8 [Mixed mode XML query processing](#)

Alan Halverson, Josef Burger, Leonidas Galanis, Ameet Kini, Rajasekar Krishnamurthy, Nagaraja Rao, Feng Tian, Stratis D. Viglas, Yuan Wang, Jeffrey F. Naughton, David J. L. September 2003 **VLDB '2003: Proceedings of the 29th international conference on Very large data bases - Volume 29**, Volume 29


Publisher: VLDB Endowment

Full text available:  Pdf (156.23 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Bibliometrics: Downloads (6 Weeks): 1, Downloads (12 Months): 22, Citation Count: 6

Querying XML documents typically involves both tree-based navigation and pattern matching similar to that used in structured information retrieval domains. In this paper we show that for good performance, a native XML query processing system should

9 [A light-weight distributed scheme for detecting ip prefix hijacks in real-time](#)

 Changxi Zheng, Lusheng Ji, Dan Pei, Jia Wang, Paul Francis August 2007 **SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications**

Publisher: ACM

Full text available:  Pdf (481.39 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 31, Downloads (12 Months): 144, Citation Count: 5

As more and more Internet IP prefix hijacking incidents are being reported, the value of hijacking detection services has become evident. Most of the current hijacking detection approaches monitor IP prefixes on the control plane and detect inconsistencies ...

Keywords: BGP, detection, hijacking, interception, routing


Also published in:

October 2007 **SIGCOMM Computer Communication Review** Volume 37 Issue 4

10 [Approximate query processing using wavelets](#)

Kaushik Chakrabarti, Minas Garofalakis, Rajeev Rastogi, Kyuseok Shim September 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 2-3

Publisher: Springer-Verlag New York, Inc.


Full text available:  Pdf (390.24 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [terms](#)

Bibliometrics: Downloads (6 Weeks): 12, Downloads (12 Months): 94, Citation Count: 22

Approximate query processing has emerged as a cost-effective approach for dealing with the huge data volumes and stringent response-time requirements of today's decision support systems (DSS). Most work in this area, however, has so far been limited ..

Keywords: Approximate query answers, Data synopses, Query processing, Wavelet decomposition

11 Bloom filter based routing for content-based publish/subscribe

 Zbigniew Jerzak, Christof Fetzer

July 2008 **DEBS '08**: Proceedings of the second international conference on Distributed event-based systems

Publisher: ACM

Full text available:  Pdf (405.55 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 40, Downloads (12 Months): 173, Citation Count: 1

Achieving expressive and efficient content-based routing in publish/subscribe systems is a difficult problem. Traditional approaches prove to be either inefficient or severely limited in their expressiveness and flexibility. We present a novel routing algorithm.

Keywords: Bloom filter, distribution, interaction, publish/subscribe

12 Keyword query cleaning

Ken Q. Pu, Xiaohui Yu

August 2008 **Proceedings of the VLDB Endowment**, Volume 1 Issue 1


Publisher: VLDB Endowment

Full text available:  Pdf (561.48 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 15, Downloads (12 Months): 74, Citation Count: 0

Unlike traditional database queries, keyword queries do not adhere to predefined schemas and are often dirty with irrelevant words from natural languages. This makes accurate and efficient keyword query processing over databases a very challenging task.

13 Estimating the selectivity of approximate string queries

 Arturas Mazeika, Michael H. Böhlen, Nick Koudas, Divyesh Srivastava

June 2007 **Transactions on Database Systems (TODS)**, Volume 32 Issue 2

Publisher: ACM

Full text available:  Pdf (364.78 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 13, Downloads (12 Months): 184, Citation Count: 2

Approximate queries on string data are important due to the prevalence of such databases and various conventions and errors in string data. We present the VSol estimator, a novel technique for estimating the selectivity of approximate string queries.


Keywords: Inverse strings, min-wise hash signatures, q-grams

14 Impact of configuration errors on DNS robustness

 Vasileios Pappas, Zhiguo Xu, Songwu Lu, Daniel Massey, Andreas Terzis, Lixia Zhang

August 2004 **SI GCOMM '04**: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications

Publisher: ACM

Full text available:  Pdf (327.13 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by terms](#)

Bibliometrics: Downloads (6 Weeks): 15, Downloads (12 Months): 124, Citation Count: 8

During the past twenty years the Domain Name System (DNS) has sustained phenomenal growth while maintaining satisfactory performance. However, the original design focused mainly on system robustness against physical failures, and neglected impact ...

Keywords: DNS, misconfigurations, resiliency

Also published in:

August 2004 **SIGCOMM Computer Communication Review** Volume 34 Issue 4

15 [Query optimization in compressed database systems](#)



Zhiyuan Chen, Johannes Gehrke, Filip Korn

June 2001 **SIGMOD '01: Proceedings of the 2001 ACM SIGMOD international conference on Management of data**

Publisher: ACM

Full text available: Pdf (262.53 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [terms](#)

Bibliometrics: Downloads (6 Weeks): 10, Downloads (12 Months): 108, Citation Count: 13

Over the last decades, improvements in CPU speed have outpaced improvements in memory and disk access rates by orders of magnitude, enabling the use of data compression techniques to improve the performance of database systems. Previous describes ...

Also published in:

June 2001 **SIGMOD Record** Volume 30 Issue 2

16 [A performance study of four index structures for set-valued attributes of low cardinality](#)

Sven Helmer, Guido Moerkotte

October 2003 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 12 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (271.43 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [terms](#)

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 44, Citation Count: 2

Abstract. The efficient retrieval of data items on set-valued attributes is an important research topic that has attracted little attention so far. We studied and modified four index structures (sequential signature files, signature trees, extendible ...

Keywords: Access methods, Database management systems, Index structures, Partition design, Set-valued attributes

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)